



(12) **EUROPEAN PATENT APPLICATION**

(43) Date of publication:
14.02.2018 Bulletin 2018/07

(51) Int Cl.:
G06F 12/0831 ^(2016.01) **G06F 12/0868** ^(2016.01)
G06F 12/1009 ^(2016.01) **G06F 12/1081** ^(2016.01)

(21) Application number: **17185833.5**

(22) Date of filing: **11.08.2017**

(84) Designated Contracting States:
AL AT BE BG CH CY CZ DE DK EE ES FI FR GB GR HR HU IE IS IT LI LT LU LV MC MK MT NL NO PL PT RO RS SE SI SK SM TR
Designated Extension States:
BA ME
Designated Validation States:
MA MD

(72) Inventors:
• **Coburn, Joel Dylan**
Mountain View, CA 94043 (US)
• **Borchers, Albert**
Mountain View, CA 94043 (US)
• **Johnson, Christopher Lyle**
Mountain View, CA 94043 (US)
• **Sprinkle, Robert S.**
Mountain View, CA 94043 (US)

(30) Priority: **12.08.2016 US 201615235495**
12.08.2016 US 201615236171

(74) Representative: **Betten & Resch**
Patent- und Rechtsanwälte PartGmbB
Maximiliansplatz 14
80333 München (DE)

(71) Applicant: **Google LLC**
Mountain View, CA 94043 (US)

(54) **HYBRID MEMORY MANAGEMENT**

(57) Methods, systems, and apparatus for receiving a request to access, from a main memory, a first cache line of a page of data; determining that the first cache line is not stored in the main memory and is stored in a secondary memory, and in response: transferring the first cache line of the page of data from the secondary memory to the main memory without transferring the entire

page of data, wherein a remaining portion of the page of data remains stored in the secondary memory; updating a page table entry associated with the page of data to point to a location of the page of data in the main memory; and transferring the remaining portion of the page of data from the secondary memory to the main memory.




FIG. 1

Description

CROSS-REFERENCE TO RELATED APPLICATION

[0001] This application claims priority to U.S. Application No. 15/235,495, entitled HYBRID MEMORY MANAGEMENT, filed on August 12, 2016. The disclosure of the foregoing application is incorporated herein by reference in its entirety.

FIELD

[0002] The present specification generally relates to memory systems.

BACKGROUND

[0003] A wide variety of memory devices can be used to maintain and store data and instructions for various computers and similar systems. In conventional computing systems, Dynamic Random Access Memory (DRAM) technology has typically been employed to operate the dynamic memory of the computer in order for an application to operate at high speeds. However, DRAM used as main memory in computer systems is no longer scaling as rapidly as in the past. As a result, DRAM storage has become a limited resource in computing environments.

[0004] A second tier of memory, such as non-volatile memory, may be used, which bridges the performance gap between DRAM and slower technologies such as NAND flash. The second tier of memory may be accessed locally over a memory or IO bus, or remotely over a high-speed network. However, applications need to explicitly manage data placement or the system must provide automatic management that transparently moves data between memory tiers. Thus, new techniques are needed for automatic management with minimal performance impact to overcome the inadequacies of existing techniques that were designed for slower storage devices.

SUMMARY

[0005] One innovative aspect of the subject matter described in this specification is embodied in systems, methods, and apparatus that include receiving a request to access, from a main memory, a first cache line of a page of data, and determining that the first cache line is not stored in the main memory and is stored in a secondary memory. In response to determining that the first cache line is not stored in the main memory and is stored in a secondary memory, transferring the first cache line of the page of data from the secondary memory to the main memory without transferring the entire page of data, wherein a remaining portion of the page of data remains stored in the secondary memory; updating a page table entry associated with the page of data to point to a location of the page of data in the main memory; and trans-

ferring the remaining portion of the page of data from the secondary memory to the main memory.

[0006] In certain implementations, a cache coherency mechanism is implemented while transferring the remaining portion of the page of data from the secondary memory to the main memory to service an access of a second cache line of the page of data.

[0007] In certain implementations, the cache coherency mechanism comprises issuing, by a cache agent, a request for ownership of cache lines in the page of data, and releasing ownership of each cache line in response to the cache line being transferred from the secondary memory to the main memory.

[0008] In certain implementations, the systems, methods, and apparatus include storing cache lines of the page of data in cache of the cache agent. In certain implementations, the systems, methods, and apparatus include storing tags and coherence state bits associated with cache lines of the page of data.

[0009] Another aspect of the subject matter described in this specification may be embodied in systems, methods, and apparatus that include, in response to a request to access the second cache line of the page of data while the remaining portion of the page of data is being transferred, recognizing a snoop request for the second cache line, and returning the second cache line in response to the snoop request once the second cache line is transferred from the secondary memory.

[0010] Yet another aspect of the subject matter described in this specification may be embodied in systems, methods, and apparatus that include, if the main memory cannot receive a page transfer, determining one of a least used pages stored in the main memory and releasing the one of the least used pages into the secondary memory, and transferring the page of data from the secondary memory to the main memory at the location of the released one of the least used pages.

[0011] Other embodiments of these aspects include corresponding systems, apparatus, and computer programs, configured to perform the actions of the methods, encoded on computer storage devices.

[0012] Particular embodiments of the subject matter described in this specification can be implemented so as to realize one or more of the following advantages. For example, the usage or access statistics for pages in memory determined by hardware may be more precise and accurate than current methods involving software and sampling techniques. Further, by transferring a particular cache line for which access is requested first, rather than the full page of data, the delay caused by transferring a page in response to a request to access the page may be reduced, which may result in faster execution of the application or process being executed.

[0013] The details of one or more embodiments of the invention are set forth in the accompanying drawings and the description below. Other features and advantages of the invention will become apparent from the description, the drawings, and the claims.

BRIEF DESCRIPTION OF THE DRAWINGS

[0014]

FIG. 1 depicts an example of a system including a memory device according to implementations of the present disclosure.

FIG. 2 depicts an example of a system including a memory device according to implementations of the present disclosure.

FIG. 3A depicts a flowchart of an example of a process for handling a page fault according to implementations of the present disclosure.

FIG. 3B depicts a flowchart of an example of a process for a write back thread according to implementations of the present disclosure.

FIG. 4 depicts a flowchart of an example of a process for determining a frequency of usage of a page with hardware and servicing a page fault using a DMA engine, according to implementations of the present disclosure.

FIG. 5 depicts a flowchart of an example of a process for determining a frequency of usage of a page with hardware and performing page migration based on the determined usage, according to implementations of the present disclosure.

FIG. 6 depicts a flowchart of an example of a process for handling a page fault by transferring a critical cache line first, according to implementations of the present disclosure.

[0015] Like reference numbers and designations in the various drawings indicate like elements.

DETAILED DESCRIPTION

[0016] New non-volatile memory technologies may be accessed locally over a memory or IO bus, or remotely over a high-speed network. The performance of these technologies has two significant implications: first, they are fast enough that software overheads to access them must be minimal; and second, they are slow enough that DRAM may typically be used as a cache. To capitalize on the advantages of these types of non-volatile memory, applications need to explicitly manage data placement or the system must provide automatic management that transparently moves data between different portions of memory.

[0017] Accordingly, embodiments of the present disclosure provide a system for high-performance automatic management of a secondary memory available either locally over a memory or IO bus, or remotely over a net-

work. The secondary memory may be computer memory that is non-volatile and persistent in nature. The secondary memory may not be directly accessed by the processor and may be slower than the primary or main memory. The main memory, also referred to as primary memory, primary storage, internal memory or first-tier memory, may be directly accessible to the CPU. As described in more detail below, hardware utilizes the cache coherence protocol and interconnect in modern processor-based systems to reduce access latencies and maintain consistency. In addition, an optimized kernel driver, for example, may provide a fast path to the second tier of memory and handle all communication with the memory management hardware. That process is advantageous compared to existing paths through the kernel for paging, which incur large costs for things like synchronization, memory management, and block IO transfers.

[0018] These features and additional features are described in more detail below.

[0019] FIG. 1 depicts an example of a system 100 including a memory device according to implementations of the present disclosure. A central processing unit (CPU) 110 may be in communication with a DRAM 120 and a memory management unit (MMU) 150. The system 100 may further include remote memory 130, which may be accessed over a network. A cache agent 140 and a home agent 145 may operate along with MMU 150 in the management of memory. The term cache agent refers to a processing unit agent that stores cache data and can initiate memory requests. In other words, the term cache agent may represent an entity which may initiate transactions into coherent memory, and which may retain copies in its own cache structure. The home agent 145 may be, for example, a memory controller. The term home agent refers to a processing unit agent that supervises or owns a region or portion of memory. In other words, the term home agent may represent an entity which services coherent transactions, supervises a portion of the coherent memory, and/or manages conflicts that may arise among different cache agents. In addition, a page table walker 160 and a translation lookaside buffer (TLB) 165 may be part of or implemented with MMU 150.

[0020] The MMU 150 is a hardware unit that may have memory references passed through it, performing the translation of virtual memory addresses to physical addresses and handling cache control. For example, the MMU 150 may use a page table as an in-memory table containing one page table entry (PTE) per page, to map virtual page numbers to physical page numbers in main memory. The translation lookaside buffer 165, as an associative cache of PTEs, may be used to avoid the necessity of accessing the main memory every time a virtual address is mapped. When a PTE prohibits access to a virtual page, for example because no physical random access memory has been allocated to that virtual page, the MMU 150 may signal a page fault to the CPU 110. Referring to FIG. 1, the system 100 may additionally include non-volatile memory 180, and may optionally in-

clude private DRAM 170.

[0021] FIG. 2 depicts an example of a system 200 including a memory device according to implementations of the present disclosure. The system 200 may include a CPU 210 and a DRAM 220. The system 200 may further include a direct memory access (DMA) engine 230 and a secondary storage 240. The CPU 210, DRAM 220, DMA engine 230, and secondary storage 240 may be in data communication via one or more busses 215, 225, or other means of data communication.

[0022] In operation, the DMA engine 230 may be used for servicing page faults, handling write backs, or page migration. The secondary storage 240 may be a memory device serving as second-tier memory, which may not be directly accessed by the CPU 210. The DMA engine 230 may be directly connected to CPU 210 or may be part of separate entity such as a cache agent. When the CPU 210 attempts to access a page from the DRAM 220, as main memory, i.e., first-tier memory or primary memory, and that page is paged out and exists in the secondary storage 240, a page fault may occur. The page fault may indicate that the operating system (OS) needs to move the page from the secondary storage 240 to the main memory, DRAM 220. According to embodiments described herein, the hardware DMA engine 230 manages the page fault, detecting the attempt to access the page that is not in main memory DRAM 220. In response, the execution of the thread stalls while the page fault is serviced. The thread may then be released to access the main memory DRAM 220 and continue execution.

[0023] In that manner, hardware may initiate the page transfer by servicing page faults using the DMA engine 230. Accordingly, there is a programmable DMA engine 230 to accelerate page transfers in servicing page faults, writing back dirty pages, and migrating pages. The DMA engine 230 may manage the servicing of multiple concurrent requests. The DMA engine 230 requests may be initiated and completed using a low overhead command and response interface over a coherent interconnect, which may be an interconnect for the shared memory. The DMA engine 230 may be configured to be a user of coherent memory, employing a memory coherence protocol, so that the DMA engine 230 may transfer data directly into the last-level cache of the processor or CPU 210.

[0024] FIG. 3A depicts a flowchart of an example of a process 300 for handling a page fault according to implementations of the present disclosure. As described above, a page fault may occur when a thread or running program accesses a memory page that is mapped into the virtual address space, but not actually loaded into main memory. The processor's MMU 150 may detect the page fault. When the page fault is detected at 310, a determination may be made at 320 as to whether there is a free page. If there is a free page, "yes", page data may be copied from the secondary storage or backing store to the free page location at 330. If there is not a free page, "no", a page may be pulled, at 360, from the

FIFO queue, which may track all the pages in memory in a queue, with the most recent arrival at the back, and the oldest arrival in front. If that page is dirty, i.e., has been modified, at 370 the system may write the page to the secondary memory or backing store. In transferring the page from the main memory to the secondary memory, at 380 the page table entry associated with the page may be invalidated, and a TLB shutdown for any entries associated with the page may be executed, e.g., causing the TLB entries to be flushed for other processors. With that page now free, page data may be copied from the secondary storage or backing store to the free page location at 330. The page tables may be updated at 340 to create a valid PTE by updating the PTE associated with the page to point to the location in the main memory of the page. Once the page fault is handled, the thread or running program may resume at 350.

[0025] FIG. 3B depicts a flowchart of an example of a process 305 for a write back thread according to implementations of the present disclosure. Upon wake up at 315, a determination may be made as to whether the number of free pages satisfy a threshold number of pages, which may be any number of pages the given architecture is configured to handle. For example, at 325, a determination may be made as to whether the number of free pages is less than or equal to a threshold. If the number of free pages does not satisfy the threshold ("no"), e.g., is greater than the threshold, the thread may sleep at 335. If, however, it is determined at 325 that the number of free pages satisfies the threshold ("yes"), the system may determine the next page in the FIFO queue at 345. If that page is dirty, i.e., has been modified, at 370 the system may write the page to the secondary memory or backing store. In transferring the page from the main memory to the secondary memory, at 380 the page table entry associated with the page may be invalidated, and a TLB shutdown for any entries associated with the page may be executed. With that page now free, a determination may once again be made as to whether the number of free pages satisfy the threshold at 325.

[0026] Page table scanning, i.e., a scan through the page table, may be performed via hardware, rather than software requiring CPU overhead, which often results in slower memory access and processing and discarding some otherwise useful cache information. In general, the frequency of access of page data, e.g., which page data is frequently accessed relative to other pages of data and which page data is infrequently accessed relative to other pages of data may be determined by scanning the page table. Each page mapped in the page table, e.g., each PTE, may have a flag or access bit that may be set whenever the page is accessed, and then cleared by the CPU after scanning the page table.

[0027] This hardware may be implemented by augmenting the page table walker or MMU, which may include one or more page table walkers, e.g., built-in hardware to read the page table and automatically load virtual-to-physical translations into the TLB. Thus, the hardware

may be part of the processor architecture, using the page table scanning mechanism in the processor. For example, the hardware may implement a routine to scan through a page table, scanning the PTEs, to determine if an access bit has been set at each PTE since the last scan. The access bit may be cleared after determining whether the access bit has been set, and then operations may proceed until the next scan of the page table. The scanning may occur periodically, e.g., with a predetermined time period between scans, or the scanning may be triggered by some external event. Each time it is determined that an access bit or flag is not set, a count may be incremented for each page, respectively.

[0028] Over time, a profile may be created from the scanning, and the profile indicates how often and/or how recently each page is accessed. For example, the hardware may include one or more counters for each page or bloom filters to maintain the usage statistics, e.g., the count of unset access bits or flags, or the results may be stored in memory, e.g., a two-level memory, to allow sorting and filtering, e.g., most and least used pages or more frequently and less frequently accessed pages. In more detail, the hardware may maintain a per-page counter to determine how recently a page was accessed, and each counter may be updated when the respective PTE is scanned. The per-page counters may be provided in on-chip SRAM for fast access. Alternatively, because the size of the secondary memory may be large, making the area cost of the counters higher, counting bloom filters may be used to maintain the access statistics about sets of pages, thereby saving area. Alternatively, the hardware may use a small amount of private DRAM or may store the counters in system DRAM.

[0029] In certain embodiments, the page table scanning to collect access statistics may be implemented with a hardware module contained in or tightly coupled with the page table walker of the processor, or may be implemented in the MMU, since the MMU may have one or more page table walkers already available to 10 devices. Accordingly, as described above, that hardware may scan the page tables and may use the access bit of each PTE to determine each the age of each page based on determining whether the access bit has been set since the last scan. For example, if the access bit is set, then the page was accessed within the last scan period; otherwise if the access bit is not set, it is has not been accessed and may be considered cold, e.g., not accessed during the last scan period. Advantageously, these usage or access statistics determined by hardware may be more precise and accurate than current methods involving software and sampling techniques.

[0030] Thus, for main memory, the number of times the access bit has not been set may be counted for each page to identify cold pages, where, for example, the higher the count for a page, the colder the page is. The count may be reset to zero when the page is accessed and the access bit is set. The page tables may be scanned periodically to detect changes to the access bit, and all ac-

cess bits may be cleared on each scan. Thus, the count may represent the age of the page, i.e., how long the page has gone un-accessed.

[0031] As described above, the DMA engine 230 may be used in servicing page faults. In certain embodiments, the access statistics may be used to determine the destination of a page transfer or to determine a policy that may be implemented to identify which pages to transfer and when to transfer those pages. For example, based on the access statistics, the pages may be ordered from most used to least used or vice versa. When the page fault is serviced using the DMA engine 230, if there are no free pages in main memory DRAM 220, a hardware-assisted paging process may release or write back one of the least used pages into the secondary memory and may use the location of that least used page to transfer the new page into main memory.

[0032] FIG. 4 depicts a flowchart of an example of a process 400 for determining a frequency of usage of a page with hardware and servicing a page fault using a DMA engine, according to implementations of the present disclosure. The process 400 includes, at 410, determining, using the processor, whether an access bit is set for each PTE of a page table based on a scan of the page table with a page table scanner. For, example, at least one page table walker of the processor may be modified to scan the page table for whether the access bit is set. For example, a thread may be added to the page table walker to scan the page table or the scanning may be done by hardware separately from the page table walker. The access bit may indicate whether a page associated with the PTE was accessed within the last scan period, e.g., between the last scan and the current scan of the page table.

[0033] At 420, a count for each page may be incremented in response to determining that the access bit is not set for the PTE associated with the page. After determining whether the access bit is set for each PTE, at 430, the processor may reset the access bit. The count for each page may be maintained by a counter for each page stored in static random access memory or counting bloom filters may maintain the access statistics about sets of pages. Alternatively, a counter for each page may be stored in dynamic random access memory, e.g., a small amount of private DRAM or in system DRAM.

[0034] The process 400 may further include, at 440, receiving a request to access a page of data from the main memory. At 450, a page fault may be initiated based on determining that the page of data, for which access is requested, is not stored in the main memory. At 460, a page fault handler may detect the page fault and may service the page fault using the DMA engine 230.

[0035] For example, if the main memory does not have a free page and cannot receive the page transfer, a particular page in main memory, e.g., one of the least used pages in main memory, may be determined based on the count for each page at 470. The DMA engine 230 may manage the page transfer and the determined one

of the pages may be released or written back into the secondary memory at 470. Further, at 470, the page of data for which access is requested may be transferred from the secondary memory to the main memory at the location of the released one of the least used pages. Alternatively, if the main memory does have a free page and can receive the page transfer, the DMA engine 230 may manage transferring of the page data from the secondary memory to the main memory at 480.

[0036] In certain embodiments, when the page fault is initiated, execution of the thread or running program may be stalled while the DMA engine 230 manages data transfer to service the page fault, as described above. The thread may be released to access the page in the main memory after the page fault is serviced.

[0037] Rather than an 10 request to access data and initiating a page fault to transfer page data from the secondary memory, in some systems, the secondary memory may be accessed directly by the processor with loads and stores, using page migration. In these systems, data may not need to be accessed at page granularity. However, it may be beneficial to take advantage of the usage or access statistics described above to migrate data automatically based on whether the data is "hot" or "cold", i.e., how recently accessed, how frequently access, or the likelihood of accessing again.

[0038] In general, utilizing features of a cache coherence architecture, in which data in the secondary memory may be associated with a home agent 145 that is responsible for a region of the physical address space, e.g., all secondary memory in the region, the home agent 145 may identify every access to the memory for which the home agent is responsible and may use that identification to maintain access statistics. The home agent 145 associated with secondary memory may maintain data about which pages are being accessed and when those pages are being accessed. Based on the frequency of access, e.g., if the frequency of access satisfies a threshold, the home agent 145 may determine that the data accessed frequently may be transferred from the secondary memory to the main memory DRAM 120.

[0039] To perform that migration, it is advantageous to determine not only which pages in the secondary memory are getting "hot", i.e., an increase in frequency of access, but also to determine which pages in main memory DRAM are getting "cold", i.e., a decrease in frequency of access. In other words, determining which pages are accessed more frequently in the secondary memory, which may be slower than the main memory, and which pages are accessed less frequently in the main memory. One process to determine which pages are accessed less frequently in the main memory is described above with reference to usage or access statistics for the main memory DRAM based on, for example, a count of the access bit being unset. The system may determine when to move data from main memory DRAM to secondary memory and when to move data from secondary memory to main memory DRAM based on the access statistics

described above for the main memory and taking advantage of the home agent tied to secondary memory to determine access statistics for the secondary memory.

[0040] Thus, the hardware may determine when to migrate pages between the main memory, e.g., DRAM 220, and the secondary memory, e.g., secondary storage 240. For example, the hardware may use a combination of cache coherence messages with the home agent and PTEs with the access bit to determine the "temperature" of data, i.e., how "hot" or frequently accessed and how "cold" or infrequently accessed certain data is, respectively. Accordingly, pages that have a rate of increasing access frequency are migrated from the secondary memory to main memory. As described above, those pages that are "warming" or being accessed more frequently may be determined by the hardware monitoring memory accesses to the secondary memory via the home agent detecting these accesses because the physical address range is owned by the home agent. Based on those determined accesses, a per-page access counter may be maintained, for example, in on-chip SRAM. Alternatively, because the size of the secondary memory may be large, making the hardware area cost of the counters higher, counting bloom filters may be used to maintain the access statistics about sets of pages, thereby saving area. Alternatively, the hardware may use a small amount of private DRAM or may store the counters in system DRAM.

[0041] In addition, as described in more detail above, pages that are cooling off or being accessed less frequently may be determined by monitoring the PTEs. For example, when the inter-access time for a page satisfies an inter-access time threshold, the system may initiate transfer of the page from main memory DRAM to the secondary memory by invalidating the PTE associated with the page, executing a TLB shutdown for any entries associated with the page, and transferring the page from main memory DRAM into the secondary memory.

[0042] FIG. 5 depicts a flowchart of an example of a process 500 for determining a frequency of usage of a page with hardware and performing page migration based on the determined usage, according to implementations of the present disclosure. The process 500 includes, at 510, determining whether an access bit is set for each PTE of a page table based on a scan of the page table with a page table scanner of the processor. The access bit may indicate whether a page associated with the PTE was accessed within the last scan period, e.g., between the last scan and the current scan of the page table. At 520, a count for each page in the main memory may be incremented in response to determining that the access bit is not set for the PTE associated with the page. Alternatively, a count for each page in the main memory may be incremented in response to determining that the access bit is set for the PTE associated with the page. After determining whether the access bit is set for each PTE, at 530, the processor may reset the access bit. The count for each page in the main memory may be maintained by a counter for each page stored in static random

access memory or counting bloom filters may maintain the access statistics about sets of pages. Alternatively, a counter for each page may be stored in dynamic random access memory, e.g., a small amount of private DRAM or in system DRAM.

[0043] The process 500 may further include, at 540, determining whether a page in a secondary memory is accessed based on a home agent receiving a request to access the page in the secondary memory. In addition, a count for each page in the secondary memory may be incremented, at 540, in response to determining that the page in the secondary memory is accessed based on the home agent receiving a request to access the page in the secondary memory. Based on that count, a page in the secondary memory may be identified, at 550, as frequently accessed. And that page identified as frequently accessed may be migrated from the secondary memory to the main memory at 550. In certain embodiments, groups of pages or portions of memory may be identified as frequently accessed and migrated from the secondary memory to the main memory.

[0044] The process 500 may further include, at 560, identifying a page in the main memory as infrequently accessed based on the count for each page in the main memory. That page identified as infrequently accessed may be migrated from the main memory to the secondary memory at 560. In certain embodiments, groups of pages or portions of memory may be identified as infrequently accessed and migrated from the main memory to the secondary memory. Identifying a page in the main memory as infrequently accessed based on the count at 560 may include determining that an inter-access time, i.e., a time between accesses, for the page satisfies an inter-access time threshold, e.g., the inter-access time for the page is greater than, or greater than or equal to, an inter-access time threshold. Thus, the page or group of pages identified as infrequently accessed may be migrated from the main memory to the secondary memory in response to determining that the inter-access time for the page or group of pages satisfies the inter-access time threshold.

[0045] The process of migrating the page identified as infrequently accessed at 560 may include invalidating the PTE or PTEs associated with the page or group of pages, executing a translation lookaside buffer shoot-down for any entries associated with the page or group of pages, and transferring the page or group of pages from the main memory to the secondary memory. The count for each page in the secondary memory may be maintained by a counter for each page stored in static random access memory or counting bloom filters may maintain the access statistics about sets of pages. Alternatively, a counter for each page may be stored in dynamic random access memory, e.g., a small amount of private DRAM or in system DRAM.

[0046] As discussed above, when, for example, an application attempts to access memory, e.g., a page, and the page is not present in the main memory, a page fault occurs. The page fault may trigger a transfer of the whole

page from the secondary memory to main memory DRAM and then the application may continue executing. However, that process may cause a delay while reading the data of the page. For example, for a 4KB-sized page, the entire 4KB may be transferred, which may cause a delay while reading the 4KB of data. According to embodiments described herein, that delay may be advantageously reduced, which may result in faster execution of the application or process being executed.

[0047] Rather than transferring the entire page, e.g., the entire 4KB in the above example, the system identifies which cache line, e.g., block of data, the application attempted to access and initially transfers only that cache line from secondary memory to main memory. With the cache line transferred for which access was requested, the application may perform its read or write of that data and continue running. After that critical cache line is transferred without transferring the entire page, the remaining portion of the page, e.g., the remaining portion of the 4KB, may be transferred in the background.

[0048] In certain embodiments, the critical cache line may be provided to the CPU, and the cache coherency mechanism or protocol may be used. This process achieves faster performance for handling page faults. For example, in operation, an application may only request to access one or two cache lines from the page, and the critical cache lines may be provided without waiting for the entire page to be transferred. The remaining portion of the page may then be transferred in the background after the application is provided the requested data, e.g., the critical cache lines, to continue its execution.

[0049] Thus, by the hardware delivering the critical cache line first, the page fault latency may be decreased. The hardware may transfer the requested cache line from the secondary memory to, for example, the processor's cache and host DRAM. Accordingly, the requested data may be provided before the rest of the cache lines in the requested page are transferred. Further, for example, a custom page fault handler, or the hardware, may update the PTE to point to the new page in host DRAM. As a result, other processors may access cache lines in the page before those cache lines are actually present in main memory.

[0050] In certain embodiments, to maintain a consistent view of memory, the cache agent 140 may execute a request-for-ownership (RFO) on the cache lines in the page to hold them exclusively while the page transfer occurs. After the hardware finishes transferring a cache line, the cache agent 140 may release ownership of that cache line so that future requests may be serviced from the host's cache or host DRAM. Further, for example, when the host attempts to access one of the other cache lines in the page while the page is being transferred, the hardware may recognize a snoop request, e.g., an inter-processor cache coherency operation, for that other cache line in accordance with the cache coherence mechanism, and may provide the other cache line in response to the snoop request once the other cache line

is transferred from the secondary memory.

[0051] In certain embodiments, the hardware may not actually store the cache lines for pages in transit in the cache agent's cache. Rather, the tags and coherence state bits for those lines may be stored, and then the data for a request from either the host DRAM or from slower memory may be forwarded. In either arrangement, the on-chip SRAM required to cache pages being transferred may be reduced.

[0052] FIG. 6 depicts a flowchart of an example of a process 600 for handling a page fault by transferring a critical cache line first, according to implementations of the present disclosure. The process 600 may include, at 610, receiving a request to access a cache line of a page of data from a main memory. At 620, a determination may be made that the requested cache line is not stored in the main memory, but rather may be stored in secondary memory. The requested cache line, e.g., data block, of the page of data from the secondary memory may be transferred to the main memory, at 630, without transferring the entire page of data. Thus, the requested cache line is provided and the remaining portion of the page of data other than the requested cache line may remain stored in the secondary memory.

[0053] The process 600 may further include, at 640, updating a PTE associated with the page of data to point to a location of the page of data in the main memory. At 650, the remaining portion of the page of data may be transferred from the secondary memory to the main memory in the background after the requested cache line has been provided in response to the request to access the cache line.

[0054] In certain embodiments, a cache coherency mechanism may be implemented, while transferring the remaining portion of the page of data from the secondary memory to the main memory at 650, to service an access of a second cache line of the page of data different than the first cache line requested. The cache coherency mechanism may include issuing, by the cache agent 140, a RFO of cache lines in the page of data, and releasing ownership of each cache line in response to the cache line being transferred from the secondary memory to the main memory. In accordance with the cache coherency mechanism, cache lines of the page of data may be stored in cache of the cache agent 140. Alternatively, the tags and coherence state bits associated with those cache lines of the page of data may be stored. Further, in response to a request to access the second cache line of the page of data while the remaining portion of the page of data is being transferred at 650, a snoop request may be recognized for the second cache line, and returning the second cache line in response to the snoop request once the second cache line is transferred from the secondary memory. If the main memory cannot receive a page transfer, e.g., there are no free pages, one of the least used pages stored in the main memory may be determined based on the usage statistics described above, and the one of the least used pages may be re-

leased or written back into the secondary memory. Accordingly, the page of data may be transferred from the secondary memory to the main memory at the location of the released one of the least used pages.

[0055] A number of implementations have been described. Nevertheless, it will be understood that various modifications may be made without departing from the spirit and scope of the disclosure. For example, various forms of the flows shown above may be used, with steps re-ordered, added, or removed.

[0056] Implementations of the invention and all of the functional operations described in this specification can be implemented in digital electronic circuitry, firmware, or hardware, including the structures disclosed in this specification and their structural equivalents, or in combinations of one or more of them. Implementations of the invention can be implemented as one or more computer program products, i.e., one or more modules of computer program instructions encoded on a computer readable medium for execution by, or to control the operation of, data processing apparatus. The computer readable medium can be a machine-readable storage device, a machine-readable storage substrate, a memory device, or a combination of one or more of them. The term "data processing apparatus" encompasses all apparatus, devices, and machines for processing data, including by way of example a programmable processor, a computer, or multiple processors or computers. The apparatus can include, in addition to hardware, code that creates an execution environment for the computer program in question, e.g., code that constitutes processor firmware, a protocol stack, a database management system, an operating system, or a combination of one or more of them.

[0057] While this disclosure contains many specifics, these should not be construed as limitations on the scope of the invention or of what may be claimed, but rather as descriptions of features specific to particular implementations of the invention. Certain features that are described in this specification in the context of separate implementations can also be implemented in combination in a single implementation. Conversely, various features that are described in the context of a single implementation can also be implemented in multiple implementations separately or in any suitable subcombination. Moreover, although features may be described above as acting in certain combinations and even initially claimed as such, one or more features from a claimed combination can in some cases be excised from the combination, and the claimed combination may be directed to a subcombination or variation of a subcombination.

[0058] Similarly, while operations are depicted in the drawings in a particular order, this should not be understood as requiring that such operations be performed in the particular order shown or in sequential order, or that all illustrated operations be performed, to achieve desirable results. In certain circumstances, multitasking and parallel processing may be advantageous. Moreover, the

separation of various system components in the implementations described above should not be understood as requiring such separation in all implementations, and it should be understood that the described program components and systems can generally be integrated together in a single software product or packaged into multiple software products.

[0059] Thus, particular implementations of the present disclosure have been described. Other implementations are within the scope of the following claims. For example, the actions recited in the claims can be performed in a different order and still achieve desirable results. A number of implementations have been described. Nevertheless, it will be understood that various modifications may be made without departing from the spirit and scope of the disclosure. For example, various forms of the flows shown above may be used, with steps re-ordered, added, or removed. Accordingly, other implementations are within the scope of the following claims.

Claims

1. A computer-implemented method comprising:

receiving a request to access, from a main memory, a first cache line of a page of data;
determining that the first cache line is not stored in the main memory and is stored in a secondary memory, and in response,
transferring the first cache line of the page of data from the secondary memory to the main memory without transferring the entire page of data, wherein a remaining portion of the page of data remains stored in the secondary memory;
updating a page table entry associated with the page of data to point to a location of the page of data in the main memory; and
transferring the remaining portion of the page of data from the secondary memory to the main memory.

2. The method of claim 1, wherein a cache coherency mechanism is implemented while transferring the remaining portion of the page of data from the secondary memory to the main memory to service an access of a second cache line of the page of data.

3. The method of claim 2, wherein the cache coherency mechanism comprises:

issuing, by a cache agent, a request for ownership of cache lines in the page of data;
releasing ownership of each cache line in response to the cache line being transferred from the secondary memory to the main memory.

4. The method of claim 3, further comprising:

storing cache lines of the page of data in cache of the cache agent.

5. The method of claim 3, further comprising:

storing tags and coherence state bits associated with cache lines of the page of data.

6. The method of one of claims 3 to 5, further comprising:

in response to a request to access the second cache line of the page of data while the remaining portion of the page of data is being transferred, recognizing a snoop request for the second cache line, and returning the second cache line in response to the snoop request once the second cache line is transferred from the secondary memory.

7. The method of one of claims 1 to 6, further comprising:

if the main memory cannot receive a page transfer, determining one of a least used pages stored in the main memory and releasing the one of the least used pages into the secondary memory, and transferring the page of data from the secondary memory to the main memory at the location of the released one of the least used pages.

8. A system comprising:

one or more processors; and
one or more storage devices storing instructions that are operable, when executed by the one or more processors, to cause the one or more processors to perform operations comprising:

receiving a request to access, from a main memory, a first cache line of a page of data;
determining that the first cache line is not stored in the main memory and is stored in a secondary memory;
transferring the first cache line of the page of data from the secondary memory to the main memory without transferring the entire page of data, wherein a remaining portion of the page of data remains stored in the secondary memory;
updating a page table entry associated with the page of data to point to a location of the page of data in the main memory;
transferring the remaining portion of the page of data from the secondary memory

to the main memory.

9. The system of claim 8, wherein a cache coherency mechanism is implemented while transferring the remaining portion of the page of data from the secondary memory to the main memory to service an access of a second cache line of the page of data. 5

10. The system of claim 9, wherein the cache coherency mechanism comprises: 10
 - issuing, by a cache agent, a request for ownership of cache lines in the page of data;
 - releasing ownership of each cache line in response to the cache line being transferred from the secondary memory to the main memory. 15

11. The system of claim 10, the operations further comprising: 20
 - storing cache lines of the page of data in cache of the cache agent.

12. The system of claim 10, the operations further comprising: 25
 - storing tags and coherence state bits associated with cache lines of the page of data.

13. The system of one of claims 10 to 12, the operations further comprising: 30
 - in response to a request to access the second cache line of the page of data while the remaining portion of the page of data is being transferred, recognizing a snoop request for the second cache line, and returning the second cache line in response to the snoop request once the second cache line is transferred from the secondary memory. 35 40

14. The system of one of claims 8 to 13, the operations further comprising:
 - if the main memory cannot receive a page transfer, determining one of a least used pages stored in the main memory and releasing the one of the least used pages into the secondary memory, and transferring the page of data from the secondary memory to the main memory at the location of the released one of the least used pages. 45 50

15. A computer-readable storage device storing instructions executable by one or more processors which, upon such execution, cause the one or more processors to perform the method of one of claims 1 to 7. 55




FIG. 1




FIG. 2




FIG. 3A




FIG. 3B




FIG. 4




FIG. 5




FIG. 6



EUROPEAN SEARCH REPORT

Application Number
EP 17 18 5833

5

10

15

20

25

30

35

40

45

50

55

DOCUMENTS CONSIDERED TO BE RELEVANT			
Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (IPC)
X	US 5 361 345 A (ENGLISH ROBERT [US]) 1 November 1994 (1994-11-01)	1,7,8, 14,15	INV. G06F12/0831
Y	* figures 1-3 *	2-5,9-12	G06F12/0868
A	* column 2, line 56 - column 4, line 52 * * column 5, line 6 - column 8, line 46 * * claims 1-4 *	6,13	G06F12/1009 G06F12/1081
Y	----- EP 0 731 944 A1 (IBM [US]) 18 September 1996 (1996-09-18)	2-5,9-12	
A	* figures 1-3b * * paragraphs [0001] - [0007], [0011] - [0026], [0028] - [0046] *	6,13	
A	----- US 2009/307459 A1 (FLEMING MATTHEW D [US] ET AL) 10 December 2009 (2009-12-10) * figures 1-5 * * paragraphs [0003], [0011], [0018] - [0040] *	7,14	

			TECHNICAL FIELDS SEARCHED (IPC)
			G06F
The present search report has been drawn up for all claims			
Place of search The Hague		Date of completion of the search 4 December 2017	Examiner Freitas, Arthur
CATEGORY OF CITED DOCUMENTS X : particularly relevant if taken alone Y : particularly relevant if combined with another document of the same category A : technological background O : non-written disclosure P : intermediate document		T : theory or principle underlying the invention E : earlier patent document, but published on, or after the filing date D : document cited in the application L : document cited for other reasons & : member of the same patent family, corresponding document	

EPO FORM 1503 03.02 (P04C01)

**ANNEX TO THE EUROPEAN SEARCH REPORT
ON EUROPEAN PATENT APPLICATION NO.**

EP 17 18 5833

5 This annex lists the patent family members relating to the patent documents cited in the above-mentioned European search report.
The members are as contained in the European Patent Office EDP file on
The European Patent Office is in no way liable for these particulars which are merely given for the purpose of information.

04-12-2017

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 5361345 A	01-11-1994	DE 4227784 A1 GB 2259795 A US 5361345 A	25-03-1993 24-03-1993 01-11-1994
EP 0731944 A1	18-09-1996	AT 210855 T DE 69524564 D1 DE 69524564 T2 EP 0731944 A1 ES 2164781 T3 JP 3280207 B2 JP H08115260 A US 5613153 A WO 9611430 A2	15-12-2001 24-01-2002 22-08-2002 18-09-1996 01-03-2002 30-04-2002 07-05-1996 18-03-1997 18-04-1996
US 2009307459 A1	10-12-2009	NONE	

REFERENCES CITED IN THE DESCRIPTION

This list of references cited by the applicant is for the reader's convenience only. It does not form part of the European patent document. Even though great care has been taken in compiling the references, errors or omissions cannot be excluded and the EPO disclaims all liability in this regard.

Patent documents cited in the description

- US 23549516 A [0001]